K-Theory for Operator Algebras
MSRI Publications
Volume **5**, second edition, 1998

CHAPTER I

INTRODUCTION TO K-THEORY

1. Survey of Topological K-Theory

This expository section is intended only as motivation and historical perspective for the theory to be developed in these notes. See [Atiyah 1967; Karoubi 1978] for a complete development of the topological theory.

K-theory is the branch of algebraic topology concerned with the study of vector bundles by algebraic means. Vector bundles have long been important in geometry and topology. The first notions of K-theory were developed by Grothendieck in his work on the Riemann–Roch theorem in algebraic geometry. K-theory as a part of algebraic topology was begun by Atiyah and Hirzebruch [1961].

1.1. Vector Bundles

Informally, a vector bundle over a base space X (which we will usually take to be a compact Hausdorff space) is formed by attaching a finite-dimensional vector space to each point of X and tying them together in an appropriate manner so that the bundle itself is a topological space. More specifically:

DEFINITION 1.1.1. A *vector bundle* over X is a topological space E, a continuous map $p : E \to X$, and a finite-dimensional vector space structure on each $E_x = p^{-1}(x)$ compatible with the induced topology, such that E is *locally trivial*: for each $x \in X$ there is a neighborhood U of x such that $E|_U = p^{-1}(U)$ is isomorphic to a trivial bundle over U.

An *isomorphism* of vector bundles E and F over X is a homeomorphism from E to F which takes E_x to F_x for each $x \in X$ and which is linear on each fiber.

A *trivial bundle* over X is a bundle of the form $X \times V$, where V is a fixed finite-dimensional vector space and p is projection onto the first coordinate (the topology is the product topology).

One can consider real vector bundles or complex vector bundles (or even quaternionic vector bundles), according to whether the vector spaces are real or complex. We will later restrict attention to complex bundles since it is the complex theory which generalizes to (complex) Banach algebras, but for much of the basic theory either kind can be considered.

The local triviality implies that the dimension of the fibers is locally constant, and hence is globally constant if X is connected (really the only interesting case).

If the dimension of each fiber is n, we say the bundle is n-dimensional. A one-dimensional bundle is sometimes called a *line bundle.*

Every X has at least one bundle of each dimension, namely the trivial bundle. Many spaces have only trivial bundles.

EXAMPLES 1.1.2. (a) The simplest example of a nontrivial (real) vector bundle is the Möbius strip, formed from $[0, 1]\times\mathbb{R}$ by identifying $(0, x)$ with $(1, -x)$. It is a vector bundle over the circle S^1.

(b) Another interesting nontrivial vector bundle is the tangent bundle TS^2 of the 2-sphere S^2. More generally, many vector bundles are naturally associated to differentiable manifolds: tangent and cotangent bundles, normal bundles associated with immersions, bundles associated with foliations, etc.

(c) There is a general "clutching" construction which yields many bundles. If $X = X_1 \cup X_2$, and if E_i is a bundle on X_i with $E_1|_Y \cong E_2|_Y$, where $Y = X_1 \cap X_2$, then E_1 and E_2 can be glued together over Y to give a vector bundle over X. (Clutching can also be done more generally.) The resulting bundle depends on which isomorphism is taken between $E_1|_Y$ and $E_2|_Y$. For example, let X be S^2, X_1 the upper hemisphere, X_2 the lower hemisphere, and E_1 and E_2 trivial complex line bundles. Then $Y \cong S^1$, and $E_i|_Y$ is a trivial complex line bundle. Let σ_n be the map which sends $(z, w) \in E_1|_Y$, for $z \in S^1$ and $w \in \mathbb{C}$, to $(z, z^n w) \in E_2|_Y$. Then the bundles on S^2 corresponding to the σ_n are all mutually nonisomorphic, and all complex line bundles arise in this manner. σ_0, of course, gives the trivial bundle. In fact, any complex vector bundle on S^n is formed in a similar way by clutching over the "equator", which is an $(n-1)$-sphere.

There is an alternate way of viewing the bundle arising from σ_1, which helps motivate the algebraic reformulation of K-theory described in 1.7: identify S^2 with $\mathbb{CP}^1$, the projective space of one-dimensional (complex) subspaces of $\mathbb{C}^2$, and set $V = \{(x, v) \in \mathbb{CP}^1 \times \mathbb{C}^2 : v \in x\}$. This is a subbundle of the trivial two-dimensional bundle. It is probably the most important bundle, and is called the *Bott bundle.*

It is not always easy to tell whether a given bundle is trivial. In cases (a) and (b) above, the nontriviality can most easily be established by looking at the sections of the bundle. A *section* of a bundle E over X is a continuous function $s : X \to E$ with $s(x) \in E_x$ for each x, i.e. a continuous choice of a vector in each fiber. A trivial bundle has many globally nonvanishing sections, for example the constant sections. Neither the Möbius strip nor the tangent bundle to S^2 has a globally nonvanishing section. (For a tangent bundle, a section corresponds to a vector field on the manifold. It is well known that S^2 has no globally nonvanishing vector fields.) We will denote the set of all sections of E by $\Gamma(E)$.

It must be emphasized that we are considering vector bundles as topological objects only. In case (b) above, the bundles have a natural differentiable structure, and it is this differentiable structure which is crucial in many of the applications of vector bundle theory to differential geometry and differential topology. We will ignore the differentiable structure completely, as it is irrelevant for K-theory. (This is actually not a serious loss, since a theorem of differential topology [Hirsch 1976, 4.3.5] says that every topological vector bundle over a differentiable manifold has an essentially unique differentiable structure compatible with the manifold.) So for us, differentiable structures on X serve only to give a way of constructing interesting topological vector bundles over X. (This is not to say that differentiable structures are irrelevant for operator algebras. On the contrary, some of the deepest and most fascinating work now being done is the noncommutative differential geometry of Connes [1994]. There are some close connections between this work and K-theory, which we will unfortunately only be able to very briefly touch on in these notes.)

1.1.3. If E is a vector bundle over X, and $\phi : Y \to X$ is a continuous function, then E can be pulled back via ϕ to a vector bundle $\phi^*(E)$ over Y. $\phi^*(E)$ can be defined formally as the fibered product $Y \times_X E = \{(y,e) \in Y \times E \mid \phi(y) = p(e)\}$. The fiber $\phi^*(E)_y$ is just $E_{\phi(y)}$. The pullback of a bundle is a bundle of the same dimension, and the pullback of a trivial bundle is trivial.

1.2. Whitney Sum

There are a number of operations which can be used to combine vector bundles over X into new ones. The most important one for our purposes is the "direct sum" or Whitney sum. Given bundles E and F over X, the Whitney sum $E \oplus F$ is formed by taking the fiberwise direct sum and tying the fibers together in a way compatible with the topologies on E and F. More precisely, if p and q are the projection maps of E and F onto X, we have

$$E \oplus F = \{(e,f) \in E \times F \mid p(e) = q(f)\}.$$

The sum of an n-dimensional bundle and an m-dimensional bundle is an $(m+n)$-dimensional bundle, and the sum of two trivial bundles is a trivial bundle. Whitney sums commute with pullbacks.

The Whitney sum makes the set of isomorphism classes of complex vector bundles over X into a commutative monoid (semigroup with identity) denoted $V_{\mathbb{C}}(X)$. The trivial bundles form a submonoid isomorphic to the additive monoid of nonnegative integers. The identity is the (class of the) 0-dimensional trivial bundle. A continuous map $\phi : Y \to X$ induces a homomorphism $\phi^* : V_{\mathbb{C}}(X) \to V_{\mathbb{C}}(Y)$ (1.1.3), so $V_{\mathbb{C}}$ gives a contravariant functor from topological spaces to abelian monoids. Similarly, the real vector bundles give a monoid $V_{\mathbb{R}}(X)$.

One can also form tensor products and exterior powers of bundles. These give additional algebraic structure to $V(X)$. Unfortunately, there is no known way of extending this additional structure to the noncommutative case, so these operations remain unique to topological K-theory.

EXAMPLES 1.2.1. (a) The Whitney sum of two Möbius strips is a two-dimensional trivial bundle. $V_{\mathbb{R}}(S^1)$ is isomorphic to $\{0\} \cup (\mathbb{N} \times \mathbb{Z}_2)$.

(b) Although TS^2 is a nontrivial 2-dimensional bundle over S^2, its sum with a trivial line bundle is a trivial 3-dimensional bundle. So $V_{\mathbb{R}}(S^2)$ is a complicated monoid which does not have cancellation, i.e. $x+z=y+z$ does not imply $x=y$. There are compact manifolds X (e.g. the 5-torus $\mathbb{T}^5$) for which $V_{\mathbb{C}}(X)$ does not have cancellation.

(c) Every complex vector bundle over S^2 is a sum of line bundles, and $V_{\mathbb{C}}(S^2) \cong \{0\} \cup (\mathbb{N} \times \mathbb{Z})$.

1.2.2. An important theorem of Swan says that if E is a vector bundle over a compact Hausdorff space X, then there is a bundle F such that $E \oplus F$ is a trivial bundle.

1.3. The Grothendieck Group

If H is an abelian semigroup, then there is a universal enveloping abelian group $G(H)$ called the Grothendieck group of H. $G(H)$ can be constructed in a number of ways. For example, $G(H)$ may be defined to be the quotient of $H \times H$ under the equivalence relation $(x_1, y_1) \sim (x_2, y_2)$ if and only if there is a z with $x_1 + y_2 + z = x_2 + y_1 + z$. $G(H)$ may be thought of as the group of (equivalence classes of) formal differences of elements of H, thinking of (x, y) as $x - y$. The prototype example of this construction is the construction of $\mathbb{Z}$ from $\mathbb{N}$. $G(H)$ may also be defined by generators and relations, with generators $\{\langle x \rangle : x \in H\}$ and relations $\{\langle x \rangle + \langle y \rangle = \langle x + y \rangle : x, y \in H\}$.

There is a canonical homomorphism from H into $G(H)$ which sends x to $[(x+x, x)]$. This homomorphism is injective if and only if H has cancellation. $G(H)$ has the universal property that any homomorphism from H into an abelian group factors through $G(H)$. G gives a covariant functor from abelian semigroups to abelian groups.

1.4. The K-Groups

DEFINITION 1.4.1. If X is a compact Hausdorff space, $K(X) = K_{\mathbb{C}}(X)$ is the Grothendieck group of $V_{\mathbb{C}}(X)$. $K_{\mathbb{R}}(X)$ is the Grothendieck group of $V_{\mathbb{R}}(X)$. So $K(X)$ may be thought of as the group of equivalence classes of formal differences of vector bundles over X.

$K_{\mathbb{C}}$ and $K_{\mathbb{R}}$ are sometimes written KU and KO respectively (U for "unitary", O for "orthogonal"). $K_{\mathbb{C}}$ and $K_{\mathbb{R}}$ are contravariant functors from compact Hausdorff spaces to abelian groups.

EXAMPLES 1.4.2. (a) If X is a one-point space or [0, 1], then every bundle is trivial; $V_{\mathbb{R}}(X) \cong V_{\mathbb{C}}(X)$ is the nonnegative integers, so $K_{\mathbb{R}}(X) \cong K_{\mathbb{C}}(X) \cong \mathbb{Z}$. The same is true for any contractible space.

(b) $K_{\mathbb{R}}(S^1) \cong \mathbb{Z} \times \mathbb{Z}_2$.

(c) $K_{\mathbb{C}}(S^2) \cong \mathbb{Z}^2$.

1.5. Locally Compact Spaces

Before developing the fundamental exact sequence of K-theory, we must extend the definition of $K(X)$ to locally compact Hausdorff spaces. The reason is that the exact sequence relates the K-theory of X to that of a closed subspace Y and the complement $X \setminus Y$; even if X is compact $X \setminus Y$ will not be in general.

The same definition of $K(X)$ makes sense if X is not compact, but turns out not to be appropriate. The correct approach is to define a relative K-group for a compact pair (X, Y), where X is compact and Y is a closed subspace of X. $K(X, Y)$ may be defined to be the Grothendieck group of the semigroup consisting of triples (E, F, α), where E and F are vector bundles over X whose restrictions to Y are isomorphic and α is a fixed isomorphism from $E|_Y$ to $F|_Y$; (E, F, α) and (E', F', α') are identified if $E \cong E'$ and $F \cong F'$ under isomorphisms whose restrictions to Y intertwine α and α'. (E, F, α) is also identified with $(E \oplus G,\ F \oplus G,\ \alpha \oplus id)$. (The semigroup operation is coordinatewise Whitney sum.) We then define $K(X) = K(X^+, +)$ for X locally compact, where X^+ is the one-point compactification of X and $+$ is the point at infinity. It is easy to see that this definition agrees with the previous one if X is compact. We extend the definition of the relative group $K(X, Y)$ to the locally compact case by $K(X, Y) = K(X^+, Y^+)$. We may also define $K_{\mathbb{R}}(X, Y)$ and $K_{\mathbb{R}}(X)$ in the same manner.

$K_{\mathbb{C}}$ and $K_{\mathbb{R}}$ then give functors from the category of locally compact Hausdorff spaces and *proper* maps to abelian groups. It is usually better to think of the functors as being defined on the equivalent category of pointed compact Hausdorff spaces.

Proposition 1.5.1. *If X is locally compact, Y is a closed subspace, and $U = X \setminus Y$, then the map q from X^+ to U^+ which is the identity on U and which sends $X^+ \setminus U$ to the point at infinity induces an isomorphism between $K(X, Y)$ and $K(U)$. So the sequence*

$$K(U) \xrightarrow{q^*} K(X) \xrightarrow{i^*} K(Y)$$

is exact in the middle ($i.e.$ $\ker i^* = \operatorname{im} q^*$.) *An analogous statement is true for $K_{\mathbb{R}}$.*

1.6. Exact Sequences

It is not true that q^* is injective and i^* surjective in general. However, this exact sequence can be put into a longer exact sequence. The extension uses higher K-groups defined by suspension:

Definition 1.6.1. If X is locally compact, the (*reduced*) *suspension* SX of X is defined to be $X \times \mathbb{R}$.

The unreduced suspension of a compact space X is the quotient of $X \times [0, 1]$ obtained by collapsing $X \times \{0\}$ and $X \times \{1\}$ to single points. The reduced suspension is the analogous construction in the category of pointed spaces: form the unreduced suspension of X^+, collapse $\{+\} \times [0,1]$ to a single point, and use this as base point.

One can also suspend a map, so suspension gives a functor from the category of locally compact spaces (or pointed compact spaces) to itself.

DEFINITION 1.6.2. Set $K^0(X) = K(X)$, $K^{-n}(X) = K(S^nX) = K(X \times \mathbb{R}^n)$ for $n > 0$. Similarly, set $K_{\mathbb{R}}^{-n}(X) = K_{\mathbb{R}}(S^nX)$. [The use of negative indices is a convention intended to exhibit K-theory as a cohomology theory. Because of Bott periodicity it is irrelevant in complex K-theory, but is necessary for real K-theory.]

The situation of 1.5.1 yields a short exact sequence $K^{-n}(U) \to K^{-n}(X) \to K^{-n}(Y)$ for each n, and similarly for $K_{\mathbb{R}}^{-n}$.

We now come to the two fundamental results of K-theory:

THEOREM 1.6.3 (LONG EXACT SEQUENCE OF K-THEORY). *Let X be locally compact, Y a closed subspace, $U = X \setminus Y$. Then there is a natural connecting homomorphism $\partial : K^{-n}(Y) \to K^{-n+1}(U)$ which makes the following long sequence exact:*

$$\cdots \xrightarrow{\partial} K^{-n}(U) \xrightarrow{q^*} K^{-n}(X) \xrightarrow{\iota^*} K^{-n}(Y) \xrightarrow{\partial} K^{-n+1}(U) \xrightarrow{q^*} \cdots \xrightarrow{\iota^*} K^0(Y)$$

and similarly for $K_{\mathbb{R}}$.

THEOREM 1.6.4 (BOTT PERIODICITY). *There is a natural isomorphism between $K(X)$ and $K^{-2}(X)$, hence between $K^{-n}(X)$ and $K^{-n-2}(X)$. ("Natural" is in the sense of category theory, i.e. a natural transformation between the functors K^{-n} and K^{-n-2}.) So the long exact sequence of complex K-theory becomes a cyclic 6-term exact sequence*

$$\begin{array}{ccccc} K^0(U) & \xrightarrow{q^*} & K^0(X) & \xrightarrow{\iota^*} & K^0(Y) \\ \partial \uparrow & & & & \downarrow \partial \\ K^{-1}(Y) & \xleftarrow{\iota^*} & K^{-1}(X) & \xleftarrow{q^*} & K^{-1}(U) \end{array}$$

Bott Periodicity is the place where complex K-theory begins to differ from real K-theory. There is also periodicity in real K-theory, but the period is 8 (i.e. $K_{\mathbb{R}}(X) \cong K_{\mathbb{R}}^{-8}(X)$), so the long exact sequence of real K-theory is a cyclic 24-term exact sequence. Bott Periodicity can be understood and proved using Clifford algebras, and the difference between the real and complex cases is reflected by the greater complexity of real Clifford algebras. (See [Karoubi 1978, III.3] for an exposition of Clifford algebras and their relationship to Bott Periodicity.)

The 6-term exact sequence is one of the primary tools which allow the K-groups of standard spaces to be computed.

THEOREM 1.6.5. *Real or complex K-theory is an extraordinary cohomology theory, i.e. it is a sequence of homotopy-invariant contravariant functors from compact spaces and compact pairs to abelian groups, with a long exact sequence* (1.6.3), *and satisfying the excision and continuity axioms* (*but not the dimension axiom*).

See [Spanier 1966] or [Taylor 1975] for an explanation of these terms.

THEOREM 1.6.6 (CHERN CHARACTER). *Let X be compact. Then there are isomorphisms*

$$\chi^0 : K^0(X) \otimes \mathbb{Q} \to \bigoplus_{n\,\text{even}} H^n(X;\mathbb{Q}),$$

$$\chi^1 : K^{-1}(X) \otimes \mathbb{Q} \to \bigoplus_{n\,\text{odd}} H^n(X;\mathbb{Q}),$$

where $H^n(X;\mathbb{Q})$ denotes the n-th ordinary (*Alexander or Čech*) *cohomology group of X with coefficients in $\mathbb{Q}$.*

So, at least rationally, $K^0(X)$ is just the direct sum of the even cohomology groups of X, and $K^{-1}(X)$ the sum of the odd ones.

1.7. Algebraic Formulation of K-Theory

We now describe a way of translating K-theory into an algebraic form which admits a generalization to Banach algebras and, to a lesser extent, to general rings.

1.7.1. Let E be a complex vector bundle over a compact space X. Then the set $\Gamma(E)$ of sections of E has a natural structure as a module over the ring (algebra) $C(X)$ of all complex-valued continuous functions on X. If E is a real vector bundle, then $\Gamma(E)$ is a module over $C_{\mathbb{R}}(X)$, the ring of real-valued continuous functions. If E is a trivial bundle of dimension n, then $\Gamma(E)$ is a free module of rank n. We have $\Gamma(E \oplus F) \cong \Gamma(E) \oplus \Gamma(F)$. So since every bundle is a direct summand of a trivial bundle by Swan's Theorem (1.2.2), $\Gamma(E)$ is always a *projective* module, a direct summand of a free module. Using the compactness of X, the local triviality of E, and the finite-dimensionality of the fibers, it is easy to see that the module $\Gamma(E)$ is *finitely generated.*

Conversely, any finitely generated projective module over $C(X)$ occurs as the module of sections of a bundle. This can be seen most easily by identifying projective modules with idempotents as follows.

The finitely generated projective modules over a unital ring R are exactly the direct summands of R^n for some n. The endomorphism ring of the free module R^n is $M_n(R)$. If V and W are R-modules with $V \oplus W \cong R^n$, then the projection of $V \oplus W$ onto $V \oplus 0$ gives an idempotent in $M_n(R)$. The idempotent so constructed depends on the choice of W and n, and on the identification of $\text{End}(V \oplus W)$ with $M_n(R)$; but it is not difficult to see that it is uniquely determined by V up to similarity. (We identify $x \in M_n(R)$ with $\text{diag}(x, 0) \in$

$M_{n+k}(R)$.) Conversely, if e is an idempotent in $M_n(R) \cong \text{End}(R^n)$, then the range of e is a finitely generated projective module.

If $R = C(X)$, we may identify $M_n(R)$ with the algebra $C(X, \mathbb{M}_n)$ of continuous functions from X to $\mathbb{M}_n$. If e is an idempotent in $C(X, \mathbb{M}_n)$, then e is a continuous function from X into the set of idempotents in $\mathbb{M}_n$. We form a bundle E by attaching the range of $e_x \subseteq \mathbb{C}^n$ to x, i.e. $E = \{(x, v) \in X \times \mathbb{C}^n \mid v \in \text{range}\, e_x\}$. (It is a somewhat nontrivial fact that a bundle defined this way is locally trivial.) Thus every idempotent, and hence every finitely generated projective module, over $C(X)$ comes from a bundle. The Bott bundle of 1.1.2(c) arises this way; the corresponding projection in $M_2(C(S^2))$, called the *Bott projection*, is easily described by identifying $\mathbb{CP}^1$ with the set of rank-one projections in $\mathbb{M}_2$.

Since E and F are isomorphic as bundles if and only if $\Gamma(E) \cong \Gamma(F)$ as modules, we get an isomorphism of the monoid $V(X)$ with the monoid of isomorphism classes of finitely generated projective modules over $C(X)$, with ordinary direct sum. Alternatively, $V(X)$ is isomorphic to the monoid of equivalence classes of idempotents in $M_\infty(C(X)) = \varinjlim M_n(C(X))$.

1.7.2. If R is any unital ring, we may define the monoid $V(R)$ of isomorphism classes of finitely generated projective R-modules, or of equivalence classes of idempotents in matrix algebras over R. (If R is noncommutative, we must specify left modules or right modules; but since the categories are equivalent, the resulting monoid is the same.) Thus we can define $K_0(R)$ to be the Grothendieck group of $V(R)$. K_0 is a *covariant* functor from unital rings to abelian groups (this is why we write K_0 rather than K^0). We have $K^0(X) = K_0(C(X))$ and $K^0_{\mathbb{R}}(X) = K_0(C_{\mathbb{R}}(X))$.

If R is nonunital, we define $K_0(R)$ to be the kernel of the homomorphism from $K_0(R^+)$ to $K_0(\mathbb{Z}) \cong \mathbb{Z}$, where R^+ is R with identity adjoined. (If R is a complex algebra, adjoining an identity is usually done by adding a copy of $\mathbb{C}$. For K-theory the results are the same.)

The basic properties of K^0 carry over to this algebraic situation. If $0 \to J \to R \to R/J \to 0$ is an exact sequence of rings, we have a short exact sequence $K_0(J) \to K_0(R) \to K_0(R/J)$ which is an exact generalization of 1.5.1.

1.7.3. Although K_0 works fine in a purely algebraic setting, difficulties arise in trying to define higher algebraic K-groups due to the lack of any reasonable notion of suspension. There is a way of defining higher algebraic K-groups (due to Bass for K_1, Milnor for K_2, and Quillen for general K_n), which is satisfactory in the sense that a long exact sequence is obtained, and the groups have identifiable algebraic significance (at least for n small). However, $K_1^{alg}(C(X))$ does not agree with $K^{-1}(X)$, and the definition of $K_n^{alg}(R)$ becomes successively more complicated and technical at each step. Algebraic K-theory has nonetheless become an important branch of ring theory. See [Rosenberg 1994] for a complete development of algebraic K-theory. The algebraic K-theory of C^*-algebras is of interest in certain contexts (cf. [Connes 1994]), but is not a well-developed theory at present. (The algebraic and topological K-theories coincide for stable C^*-algebras: see [Rosenberg 1997] for an account of this and related results about the algebraic K-theory of C*-algebras.)

1.7.4. If A is a Banach algebra, there is a natural notion of suspension: SA is the Banach algebra of all continuous functions from $\mathbb{R}$ to A which vanish at infinity. We have $S(C_0(X)) \cong C_0(SX)$ if X is a locally compact Hausdorff space. We then define $K_n(A) = K_0(S^n A)$. Remarkably, all of the results of topological K-theory described in this section carry over to the Banach algebra case, notably the long exact sequence and Bott Periodicity (hence the cyclic 6-term exact sequence). Not only do the results hold true, but even the proofs, when expressed in Banach algebra language, carry over verbatim in almost all cases. In fact, the Banach algebra approach is probably the most elegant and natural way to develop topological K-theory. (This approach to topological K-theory actually appeared before Banach algebra K-theory was developed, due primarily to work of Atiyah, Karoubi, Swan, and Wood.)

1.7.5. Not everything in topological K-theory carries over to the noncommutative case, however. The tensor product of vector bundles defines a multiplication on $K^0(X)$ making it into a ring. (Actually one gets a graded ring structure on $K^*(X)$.) This ring structure can be extended to $K_0(R)$ whenever R is a commutative ring, using tensor products of modules; but there is no obvious way of putting a ring structure on $K_0(R)$ for noncommutative R. Similarly, the exterior power operations on modules have no noncommutative analog.

2. Overview of Operator K-Theory

In this section, we will give an overview of the topics to be covered in these notes. The point of view taken here is considerably different than that of Section 1, and is much more in keeping with the traditional ideas of the theory of operator algebras. No knowledge of topological K-theory is assumed; however, it is beneficial to have some understanding of the material of Section 1, particularly 1.7, to appreciate how the ideas developed.

2.1. Noncommutative Topology

The theory we will develop is the heart of the subject of *noncommutative topology*, which is the process of taking a concept from topology, rephrasing it using the (contravariant) equivalence between the category of locally compact Hausdorff spaces and the category of commutative C^*-algebras, and extending the concept in a meaningful way to the category of all C^*-algebras (or some suitable subcategory). The goal of noncommutative topology is to bring ideas, techniques, and results from topology into the study of operator algebras, and vice versa; both areas have already richly benefited from this process.

One of the motivations for developing the theory of noncommutative topology (although not the only one) is that in many instances in ordinary topology the natural object of study is a "singular space" which cannot be defined and studied in purely topological terms. Good examples are the orbit space of a group action or the leaf space of a foliation. Although the singular space X may not really exist topologically, there is often a (noncommutative) C^*-algebra which plays the role of $C_0(X)$ in an appropriate sense. Most of the applications of noncommutative topology to ordinary topology and geometry exploit this point of view.

Most of the noncommutative topology done so far has been noncommutative algebraic topology, the process of extending the functors of ordinary algebraic topology, regarded as functors from commutative C^*-algebras to abelian groups, to more general C^*-algebras. There has been little success so far in extending the standard homotopy, cohomotopy, homology, or cohomology functors (perhaps with good reason—see 22.4.2 and 22.4.3); but the functors of complex K-theory extend very nicely. We will describe these functors below purely in operator algebra terms, referring to Section 1 for the connections with topology.

2.2. The K_0-Functor

The first functor we will consider is the K_0-functor. The goal here is to define a group-valued "universal dimension function," a function D from the projections of a C^*-algebra A to an abelian group G, with the properties that $D(p) = D(q)$ whenever $p \sim q$ and $D(p+q) = D(p) + D(q)$ whenever $p \perp q$, such that any function from the projections of A to an abelian group with these properties factors through D. The prototype is the Murray–von Neumann comparison theory for finite factors, which may be regarded as the K_0-theory of factors. In this theory, the dimension function takes real-number values and completely determines equivalence of projections. In the II_1 case, the dimension group is $\mathbb{R}$, since the values of the dimension function fill up an entire interval in $\mathbb{R}$, and all of $\mathbb{R}_+$ if matrix algebras are considered.

In other cases, however, the group of real numbers is either too large or inappropriate to serve as the range group of the universal dimension function. For example, for $\mathbb{C}$, $\mathbb{M}_n$, or $\mathbb{K}$, the appropriate range group is $\mathbb{Z}$. For the CAR algebra, the proper range group is the dyadic rationals. For a direct sum of two II_1 factors, the group should be $\mathbb{R}^2$. And for an infinite factor, the requirement that the function be group-valued forces the range group to be $\{0\}$.

For any C^*-algebra A, we will define an abelian group $K_0(A)$ which is the appropriate range group for the universal dimension function on A (and on all matrix algebras over A). For A unital, the elements of $K_0(A)$ are formal differences of equivalence classes of projections in matrix algebras over A. (If A is nonunital, the proper definition is less obvious.) Matrix algebras over A must be considered, and the equivalence relation on projections "stabilized", in order to get a group structure on $K_0(A)$. The definition of K_0 is functorial, i.e. any homomorphism $\phi : A \to B$ induces a homomorphism $\phi_* : K_0(A) \to K_0(B)$. In many cases, the group $K_0(A)$ has a natural partial ordering which determines comparability of projections in A; in the particularly nice case of AF algebras, this